

DOE Centers of Excellence Performance Portability Meeting

Overview and Kickoff

Rob Neely, LLNL

April 19, 2016
Glendale, AZ



LLNL-PRES-690092

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

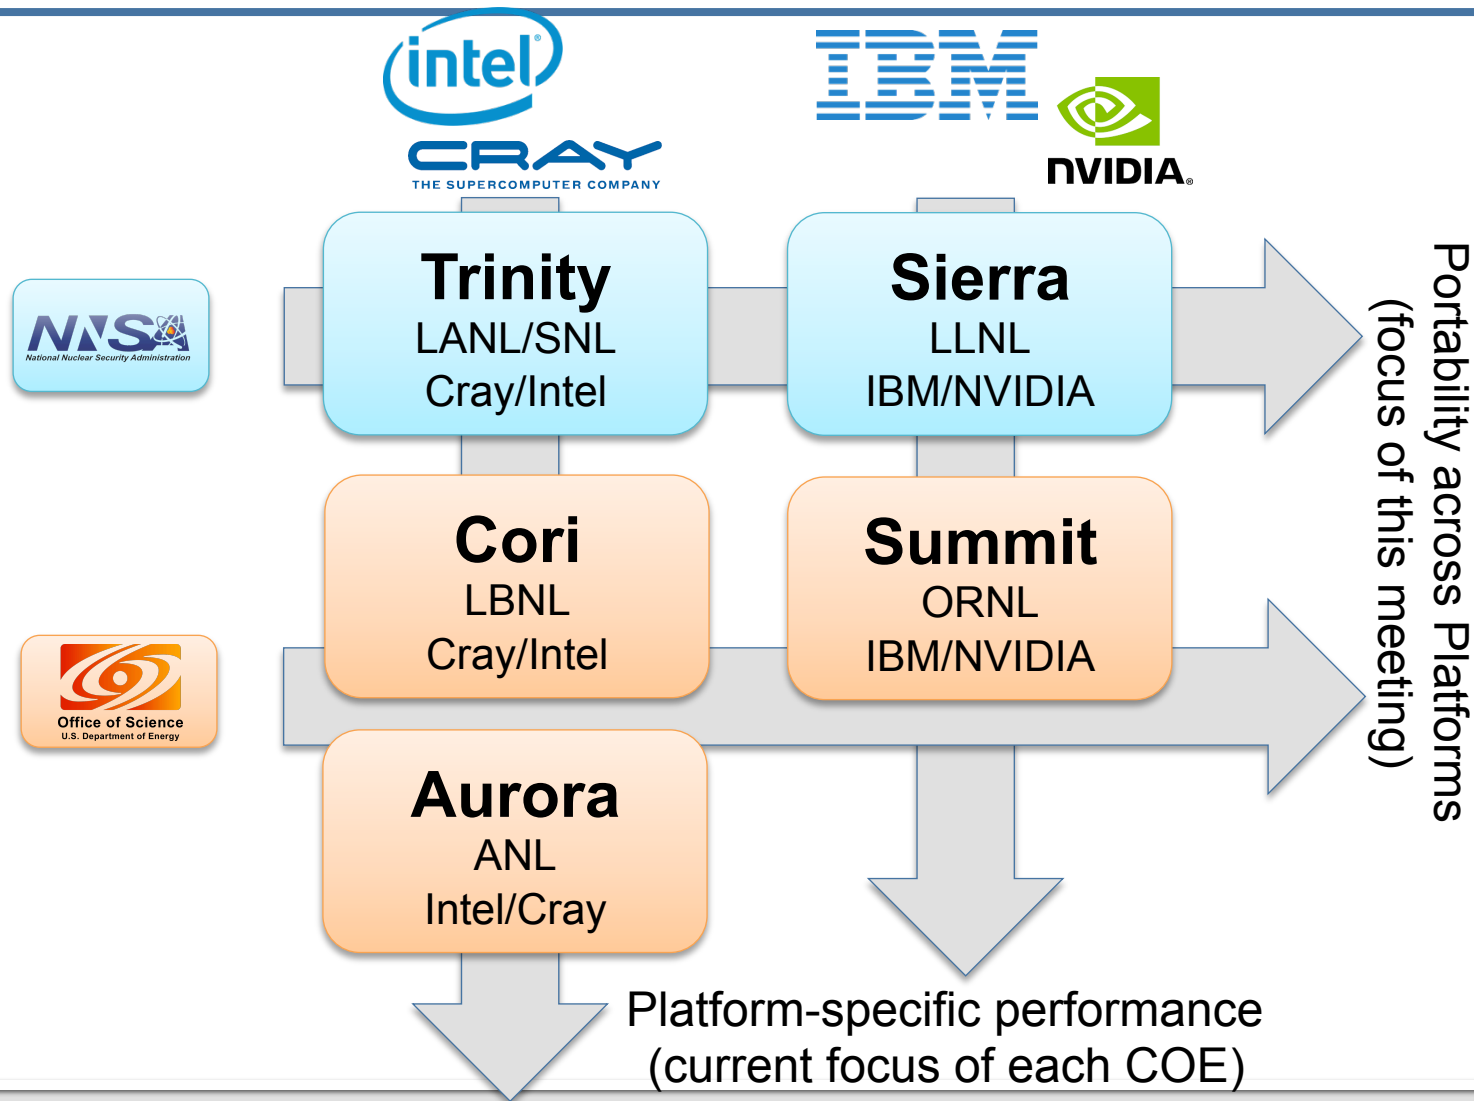
 **Lawrence Livermore
National Laboratory**

Welcome! And why are we all here?

- Center of Excellence (COE)
 - *Wikipedia: A team, a shared facility or an entity that provides leadership, best practices, research, support and/or training for a focus area.*
- COE's are now considered themselves a best practice for our large HPC procurements
 - Vendor participation provides a valuable addition to DOE application preparation
- But ...
 - There's no rulebook written for how best to utilize these partnerships
 - Cross-COE participation mostly ad-hoc due to the focused nature on a particular platform
- And...
 - Applications in DOE generally must be able to run across the Leadership Computing Facilities or ASC classified systems
 - DOE spends over \$1B dollars investing in HPC platforms every decade, and probably even more investing in application development
 - This next generation of machines are... disruptive
- So...
 - We needed a forum to raise the discussion up a notch to help bridge the excellent work going on within each COE



We're here to “diagonalize the matrix”



So what is performance portability?

- For purposes of this meeting, I propose:
 - *The ability to run an application with acceptable* performance across KNL and GPU based systems with a single version of source code*
- Both terms are subjective
 - Portable
 - From an application perspective, means not having to maintain multiple versions of algorithms tuned to different architectures
 - Performance
 - Ideally performance would be as high as is achievable using platform-specific techniques. In reality, many code teams will give up some performance for portability
- Just using portable standards (e.g. MPI+OpenMP4) does not guarantee performance portability



Goals of this meeting

- Share experience between the COEs
- Make connections between application teams working similar ideas or algorithms
- Address the numerous challenges of performance portability
- Engage vendors in helping us succeed in our performance portability goals

What are *your* goals?



What to expect this week

Tue/Wed: Lunch and Dinner on your own
Thu: Lunch and Dinner provided

- In one phrase: Breadth (not Depth)
 - Lots of shorter talks
 - Breakout sessions to identify major issues

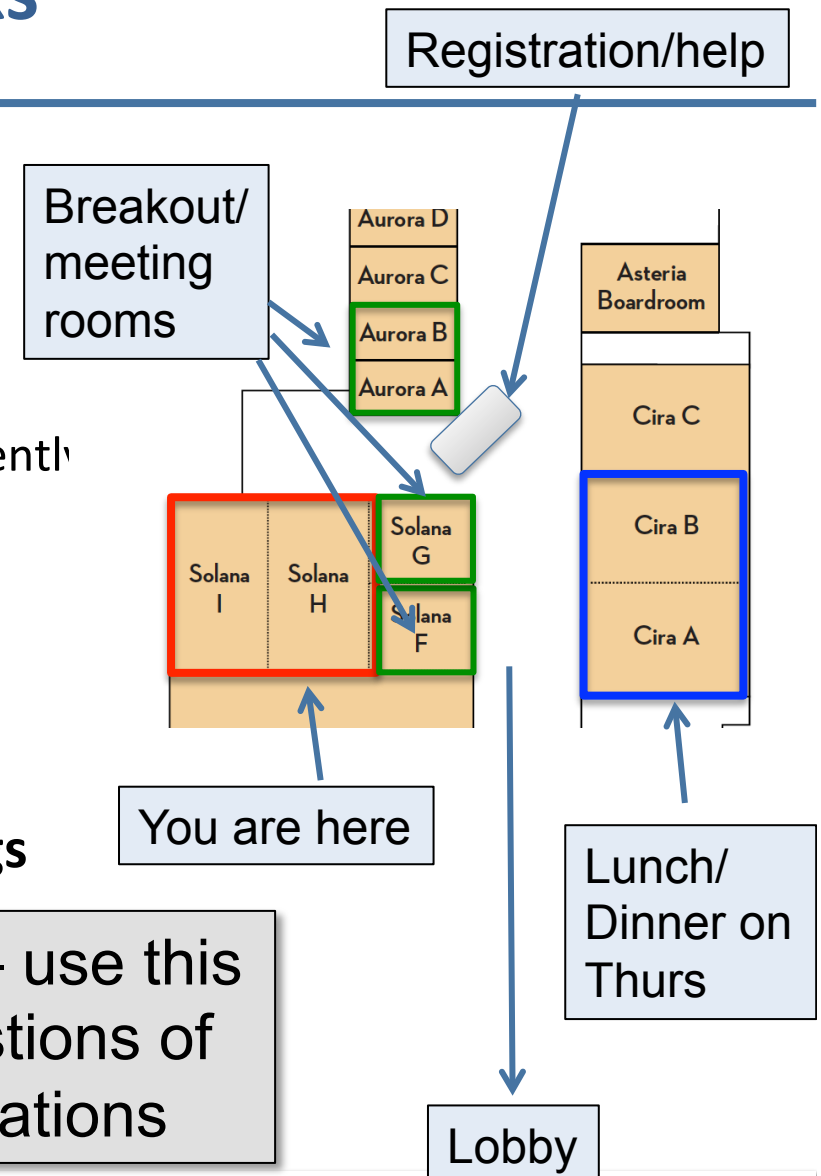
Tuesday	Wednesday	Thursday
Overviews of the five COEs Recap of HPCOR workshop	Managing the memory hierarchy	Tools & compilers
Exascale Computing Project Applications Development Focus Area	Application experiences using performance portable abstractions	IO / burst buffers
Short NDA talks from Intel and NVIDIA	Breakout sessions: <ul style="list-style-type: none">• Perf-port abstractions• Managing the mem hierarchy	Domain specific languages
Application talks – optimizations and algorithmic changes for next-gen platforms	OpenMP experiences and futures	Breakout sessions: <ul style="list-style-type: none">• OpenMP futures• Tools/compilers/system software
Performance-portable abstractions	Intel NDA session (dinner provided)	Wrapup Hosted Dinner



Breakout sessions and Breaks

- Rooms: Aurora AB, Salon F, Salon G
- Moderators will guide discussion
- Split into 4 groups (self-select).
 - Two groups covering each topic independently
- Outbriefs of discussions will follow
- Attendees are welcome to suggest alternative topics
- **Feel free to use rooms for side meetings**

Breaks, lunches, dinners – use this time to network, ask questions of speakers, build collaborations



So many talks, so little time

- Our call for abstracts exceeded expectations
 - Both in number and quality.
 - The steering committee had a difficult time selecting – all were accepted
- Thus, the time limits on talks will be strictly enforced
 - A countdown timer will inform speakers of their limit. Timer will start immediately upon transition
- Please try to hold questions until the end of the talk, and then only if extra time allows
 - Speakers: let audience know if you prefer questions during your talk
- Catch speakers during breaks



NDA material should not be discussed except in designated sessions

- Thanks to the vendors, we agreed early on to keep the meeting as open as possible
- Most everyone from the DOE labs or their affiliates are covered under NDA
 - If you are not, PLEASE do not attempt to join NDA discussions for which you are not covered
- Everyone (both speakers and participants) should have read and will adhere to our “ground rules”
 - Talks and discussions must **refrain from discussing information held under non-disclosure agreements**. Contact your steering committee representative (below) if you need specific guidance.
 - In the spirit of the meeting, talks and discussions should **address general challenges to the goal of performance portability** and approaches that might be applied to overcome those challenges, rather than identifying and comparing state-of-play at a particular point in time.
 - Talks and discussions should **not compare performance across specific platforms**. Talks and discussions can address performance improvements on a given platform due to programming approaches or can address performance achieved relative to a theoretical performance model.
 - The focus **of talks and discussions should be on portable, non-vendor-specific solutions** as seen from the application developer perspective (that is, abstractions that hide vendor-specific solutions are acceptable). It is expected that a particular focus of the meeting will be to address possible evolutions of current standards (for example, OpenMP and C++) to better support performance portability.
 - **Projections to future machines should not be presented.**
 - Talks and discussions must be **unclassified and non-sensitive** in nature.
 - Speakers and participants (both labs and vendors) should **recognize that DOE will have multiple target platforms as part of their national strategy and join the discussion in the spirit of cooperation**. All COEs are working toward the goal of making these platforms the most useful and high performance they can be without the threat of "vendor lock-in."



What to expect after the meeting concludes

- This meeting is meant to be a catalyst for future work and collaborations
 - Use this time to both learn, and build collaborations
- (Most) talks will be made available along with notes on the meeting web site (<https://asc.llnl.gov/DOE-COE-Mtg-2016>)
- An informal report will be generated to capture next steps
- Suggestions for followon work or future Multi-COE meetings should be made to anyone on the steering committee



Thanks to the Steering Committee

- This group was assembled last fall, and are what made this meeting possible

James Reinders.....	Intel/Trinity-Cori
Mike Glass.....	SNL/Trinity
Rebecca Hartman-Baker.....	LBNL/Cori
John Levesque.....	Cray/Trinity-Cori
Hai Ah Nam.....	LANL/Trinity
Rob Neely.....	LLNL/Sierra
Jim Sexton.....	IBM/Sierra-Summit
Tjerk Straatsma.....	ORNL/Summit
Tim Williams.....	ANL/Aurora
Cyril Zeller.....	NVIDIA/Sierra-Summit



Give thanks the support crew when you see them

- Lori McDowell, LLNL – primary organizer
- Ashley Wilkins, LLNL – on-site help
- Andrea Baron, LLNL – web site
- Tina Macaluso, Emily Simpson, ASC HQ - notetaking



Finally, big thanks to LLNL and the vendors

- LLNL ASC Program for meeting space,
- Vendors: Opening reception, Thursday's lunch and dinner provided with the their generous assistance
- Thanks especially to
 - Jay Gould, Cray
 - Greg Branch, NVIDIA
 - Liza Gabrielson, Intel
 - Jim Sexton and Kathryn O'brien, IBM



But it's a dry heat...

Phoenix Climate Graph - Arizona Climate Chart

High Temp in Phoenix this week

Average April High Temp in Phoenix: 85°

